

Exhibit 2 to the Complaint

U.S. Patent No. US 8,521,766 v. Verizon

Exhibit 2 to the Complaint

1. Claim Chart

Claim		Analysis
[1.P] method, comprising	A	<p>Verizon (“Company”) performs and/or induces others to perform a method.</p> <p>This element is infringed literally, or in the alternative, under the doctrine of equivalents.</p> <p>For example, Company offers to sell and induce their customers to use the infringing product iPhone 15 Pro¹ (used herein as an exemplary product), by at least teaching the step of the claimed method.</p> <p>For example, the iPhone 15 Pro comprises Siri, an intelligent voice assistant that receives voice commands from a user and retrieves information related to the voice command.</p>

¹ Company also provides iPhones including but not limited to iPhone 15, iPhone 14, iPhone SE (3rd generation), iPhone 13, iPhone 12, iPhone 11, iPhone XR, iPhone XS and all their variants, Google smartphones including but not limited to Pixel 8a, Pixel 8 Pro, Pixel 8, Pixel Fold, Pixel 7a, Pixel 7, Pixel 6a, Pixel 6, Pixel 6 Pro, and all their variants; Samsung smartphones including but not limited to Galaxy S24 Ultra, Galaxy S23 FE, Galaxy S24+, Galaxy A15 5G, Galaxy S24, Galaxy S23, Galaxy A23 5G UW, Galaxy Z Fold5, Galaxy Z Flip5, Galaxy A54 5G, Galaxy S23 Ultra, Galaxy S23+, Galaxy S21 FE 5G, Galaxy A42 5G, Galaxy S21 5G, Galaxy S21 Ultra, Galaxy S21 FE, Galaxy Note 20, Galaxy S22, Galaxy Note20 Ultra 5G, Galaxy S22+ Galaxy S20 Ultra and all their variants, that infringe in a similar way as the above exemplary product. The presently charted exemplary infringing product is representative for illustrating infringement of the other devices.

Exhibit 2 to the Complaint

✓

Mobile Home Internet Shop Deals

Search Verizon

Sign In

Free 2-day shipping on all device orders for new and existing customers. Limited time offer.

Home / Smartphones / Apple

Shop Apple smartphones

All Free phones Samsung Apple Google Motorola Kyocera

Lowest price with trade-in offer

31 results

Sort by: Featured

Apple Clear all

Filter

Brand (1)


OS

Offers

In-store pickup

Monthly price

Save \$829.99. Online only.




Apple iPhone 15

Starts at \$0/mo
~~\$22.06/mo~~ Details
for 36 months, 0% APR
Retail price: \$829.99

☐ Compare

♡

Save \$839.99. Online only.




Apple iPhone 15 Pro Max

Starts at \$10.00/mo
~~\$32.32/mo~~ Details
for 36 months, 0% APR
Retail price: \$1199.99

☐ Compare

291K ♡



Get iPhone 14 Plus, on us.
And you don't even have to
change your number.

Online only. With new line on
myPlan. Details

Buy >

3

Exhibit 2 to the Complaint

Source: <https://www.verizon.com/smartphones/apple/>

Use Siri on all your Apple devices

Use Siri to help you with the things you need to find, know or do every day. Use your voice or press a button to get Siri's attention, then say what you need. Locate your Apple device below to find out how to use Siri.

Source: <https://support.apple.com/en-us/105020>

Exhibit 2 to the Complaint

	<div><p>The diagram illustrates the architecture of Siri's voice trigger system. It shows the flow from a user's voice input to the final response. The process starts with the user saying "Siri, how's the weather?". This input goes through an "AOP or AP VoiceTrigger Detector", then an "AP VoiceTrigger Checker", and finally a "SpeakerID Trigger". The "SpeakerID Trigger" feeds into the "Siri Directed Speech Detection" block, which contains four sub-components: "SpeakerID Trigger + Payload", "Acoustic False Trigger Mitigator", "Out-of-Domain Language Detector Text-based", and "Lattice RNN ASR lattice". The output of this block goes to a "Neural Combiner", which then feeds into "Siri Natural Language Understanding". The output of NLU goes to "Text to Speech (TTS)", which finally outputs the response "It's currently partly cloudy.".</p></div> <p>Source: https://machinelearning.apple.com/research/voice-trigger#figure1</p> <p>Further, to the extent this element is performed at least in part by Company's software source code, Plaintiff shall supplement these contentions pursuant to production of such source code by the Company.</p>
[1.1] receiving an	<p>Company performs and/or induces others to perform a step of receiving an information request.</p> <p>This element is infringed literally, or in the alternative, under the doctrine of equivalents.</p>

Exhibit 2 to the Complaint

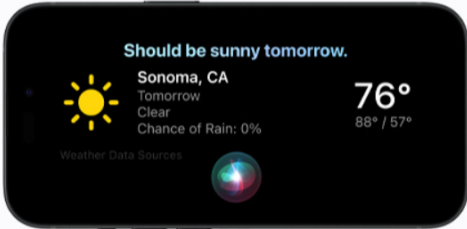
information request	<p>For example, Siri receives an input voice command (“information request”) given by a user. The command comprises a trigger phrase and a subsequent utterance, the trigger phrase being ‘Siri’ or ‘Hey Siri’.</p> <div data-bbox="411 354 1297 435"> <p>Siri, what will the weather be like tomorrow?</p> <p>information request</p> </div>  <p>Source: https://www.apple.com/siri/ (annotated)</p> <p>In this article, we will discuss how Apple has designed a high-accuracy, privacy-centric, power-efficient, on-device voice trigger system with multiple stages to enable natural voice-driven interactions with Apple devices. The voice trigger system supports several Apple device categories like iPhone, iPad, HomePod, AirPods, Mac, Apple Watch, and Apple Vision Pro. Apple devices simultaneously support two keywords for voice trigger detection: “Hey Siri” and “Siri.”</p> <p>Source: https://machinelearning.apple.com/research/voice-trigger</p> <p>Further, to the extent this element is performed at least in part by Company's software source code, Plaintiff shall supplement these contentions pursuant to production of such source code by the Company.</p>
[1.2] decoding the	Company performs and/or induces others to perform a step of decoding the information request.

Exhibit 2 to the Complaint

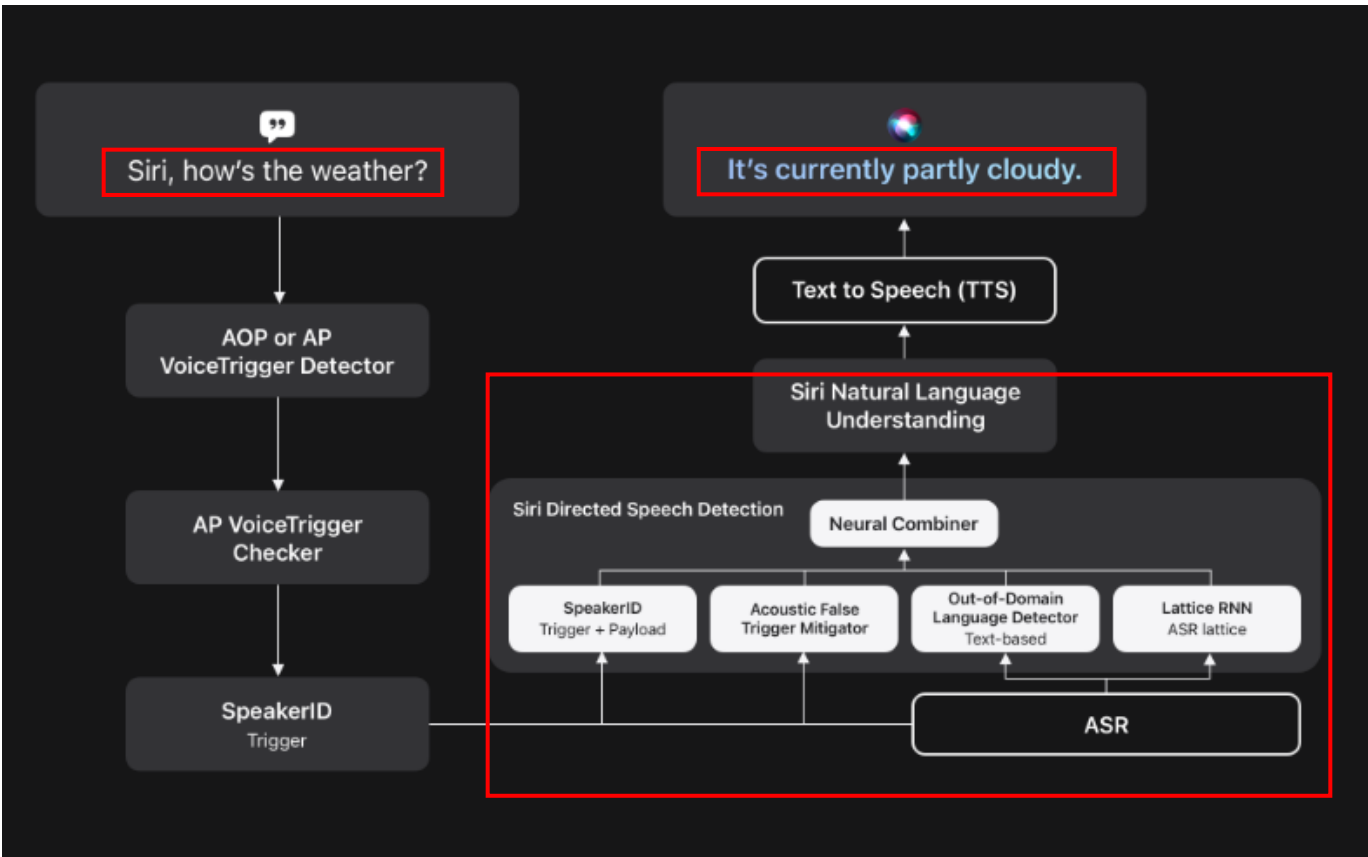
<p>information request;</p>	<p>This element is infringed literally, or in the alternative, under the doctrine of equivalents.</p> <p>For example, Siri uses automatic speech recognition and natural language processing to process the information request received from the user. It checks whether the user command is directed towards Siri or not and then, identifies the intent of the command using a Siri Directed Speech Detection (SDSD) system. The SDSD comprises various False Trigger Mitigation (FTM) systems such as Acoustic FTM, Out-of-domain Language Detector, and Lattice RNN which decode the input command and convert it into the intent.</p> 
-----------------------------	---

Exhibit 2 to the Complaint

Source: <https://machinelearning.apple.com/research/voice-trigger#figure1>

Being able to use Siri without pressing buttons is particularly useful when hands are busy, such as when cooking or driving, or when using the Apple Watch. As Figure 1 shows, the whole system has several parts. Most of the implementation of Siri is "in the Cloud", including the main automatic speech recognition, the natural language interpretation and the various information services. There are also servers that can provide updates to the acoustic models used by the detector. This article concentrates on the part that runs on your local device, such as an iPhone or Apple Watch. In particular, it focusses on the detector: a specialized speech recognizer which is always listening just for its wake-up phrase (on a recent iPhone with the "Hey Siri" feature enabled).

Source: <https://machinelearning.apple.com/research/hey-siri>

When a voice trigger detection mechanism detects a trigger, the system starts processing user audio using a full-blown ASR system. A dedicated algorithm determines the end-of-speech event, at which point we obtain the ASR output and the decoding lattice. We use word-aligned lattices such that each arc corresponds to

Source: <https://machinelearning.apple.com/research/voice-trigger>

model scores, text, etc. NLU signals are comprised of domain classification features such as domain categories, domain scores, sequence labels of the user request transcription, etc. An intent is a combination of ASR and NLU signals. We refer to these signals as *understanding signals* decoded by ASR and NLU sub-systems. Every intent is encoded into a vector space and this process is described in Section 4.1. Our task is to produce a ranked list of intents using information-state in addition to understanding signals to choose the best response.

Source: <https://arxiv.org/pdf/2005.00119.pdf>, Page 2

Exhibit 2 to the Complaint

	Further, to the extent this element is performed at least in part by Company's software source code, Plaintiff shall supplement these contentions pursuant to production of such source code by the Company.
[1.3] discovering information using the decoded information request;	<p>Company performs and/or induces others to perform a step of discovering information using the decoded information request.</p> <p>This element is infringed literally, or in the alternative, under the doctrine of equivalents.</p> <p>For example, after the intents (“decoded information request”) are determined, the relevant information is retrieved based on the intents.</p>

Exhibit 2 to the Complaint

SiriKit provides the following intents.	
Domain (link to developer guidance)	Intents
VoIP Calling	Initiate calls.
Workouts	Start, pause, resume, end, and cancel workouts.
Lists and Notes	Create notes.
	Search for notes.
	Create reminders based on a date, time, or location.
Media	Search for and play media content, such as video, music, audiobooks, and podcasts.
	Like or dislike items.
	Add items to a library or playlist.

Source: <https://developer.apple.com/design/human-interface-guidelines/siri> (annotated)

discovering
information

Exhibit 2 to the Complaint

Being able to use Siri without pressing buttons is particularly useful when hands are busy, such as when cooking or driving, or when using the Apple Watch. As Figure 1 shows, the whole system has several parts. Most of the implementation of Siri is "in the Cloud", including the main automatic speech recognition, the natural language interpretation and the various information services. There are also servers that can provide updates to the acoustic models used by the detector. This article concentrates on the part that runs on your local device, such as an iPhone or Apple Watch. In particular, it focusses on the detector: a specialized speech recognizer which is always listening just for its wake-up phrase (on a recent iPhone with the "Hey Siri" feature enabled).

Source: <https://machinelearning.apple.com/research/hey-siri>

Exhibit 2 to the Complaint

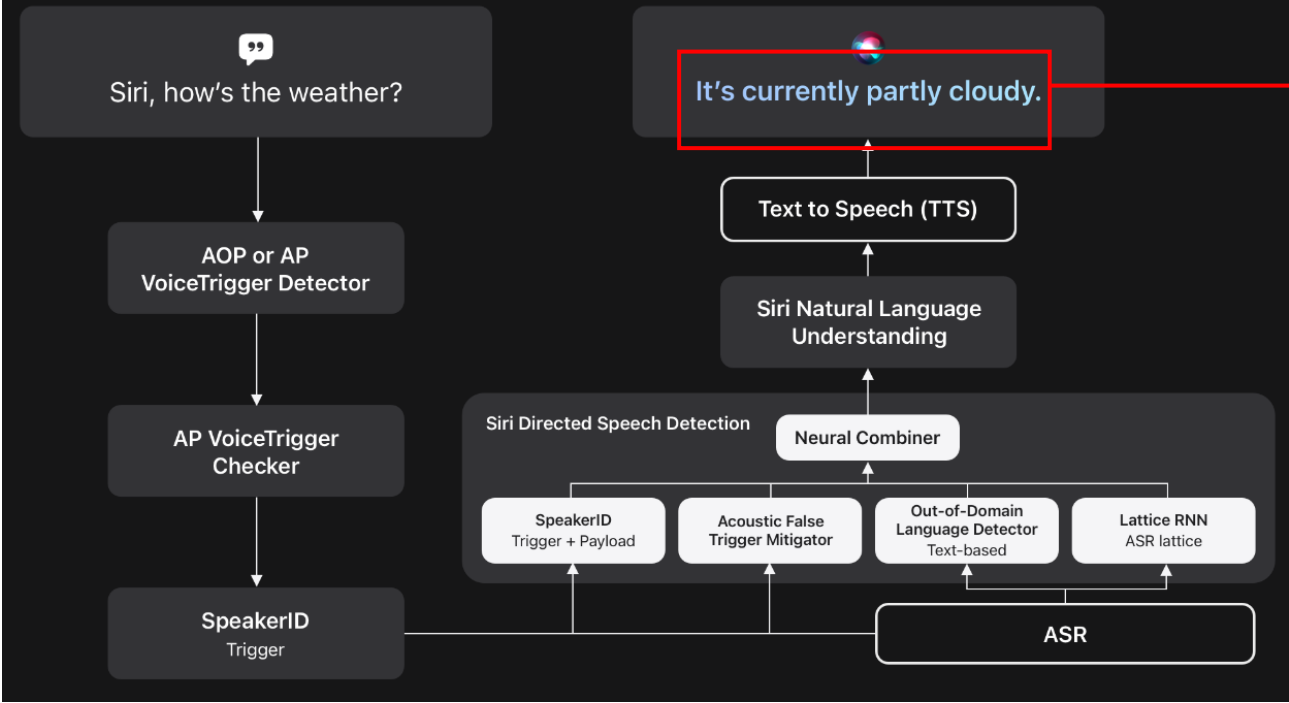
	 <p>Source: https://machinelearning.apple.com/research/voice-trigger (annotated)</p> <p>Further, to the extent this element is performed at least in part by Company's software source code, Plaintiff shall supplement these contentions pursuant to production of such source code by the Company.</p>
[1.4] preparing, using one or more processing devices instructions	<p>Company performs and/or induces others to perform a step of preparing, using one or more processing devices instructions for accessing the information.</p> <p>This element is infringed literally, or in the alternative, under the doctrine of equivalents.</p> <p>For example, after converting audio requests to the intents, Siri provides conversational flow for system intents such that the information is retrieved. These flows help the app to fulfill the user request according to the domain of the intents.</p>

Exhibit 2 to the Complaint

<p>for accessing the information, the instructions including:</p>	<p>Therefore, it would be apparent to a person having ordinary skill in the art that Siri prepares instructions for accessing the information using one or more processing devices.</p> <div data-bbox="411 342 1199 662"> <p>A closer look at intents</p> <p>When people use Siri to ask questions and perform actions, Siri does the language processing and semantic analysis needed to <u>turn their requests into intents for your app to handle</u>. The exception is the personal phrase that people create to run a shortcut: When people speak the exact phrase, Siri recognizes it without doing additional processing or analysis.</p> </div> <p>Source: https://developer.apple.com/design/human-interface-guidelines/siri/</p> <div data-bbox="411 761 1176 1050"> <p>System intents</p> <p>SiriKit defines a large number of system intents that represent common tasks people do, such as playing music, sending messages to friends, and managing notes. For system intents, Siri defines the conversational flow, while your app provides the data to complete the interaction.</p> </div> <p>Source: https://developer.apple.com/design/human-interface-guidelines/siri#System-intents</p>
---	--

Exhibit 2 to the Complaint

SiriKit provides the following intents.

Domain (link to developer guidance)	Intents
VoIP Calling	Initiate calls.
Workouts	Start, pause, resume, end, and cancel workouts.
Lists and Notes	Create notes.
	Search for notes.
	Create reminders based on a date, time, or location.
Media	Search for and play media content, such as video, music, audiobooks, and podcasts.
	Like or dislike items.
	Add items to a library or playlist.

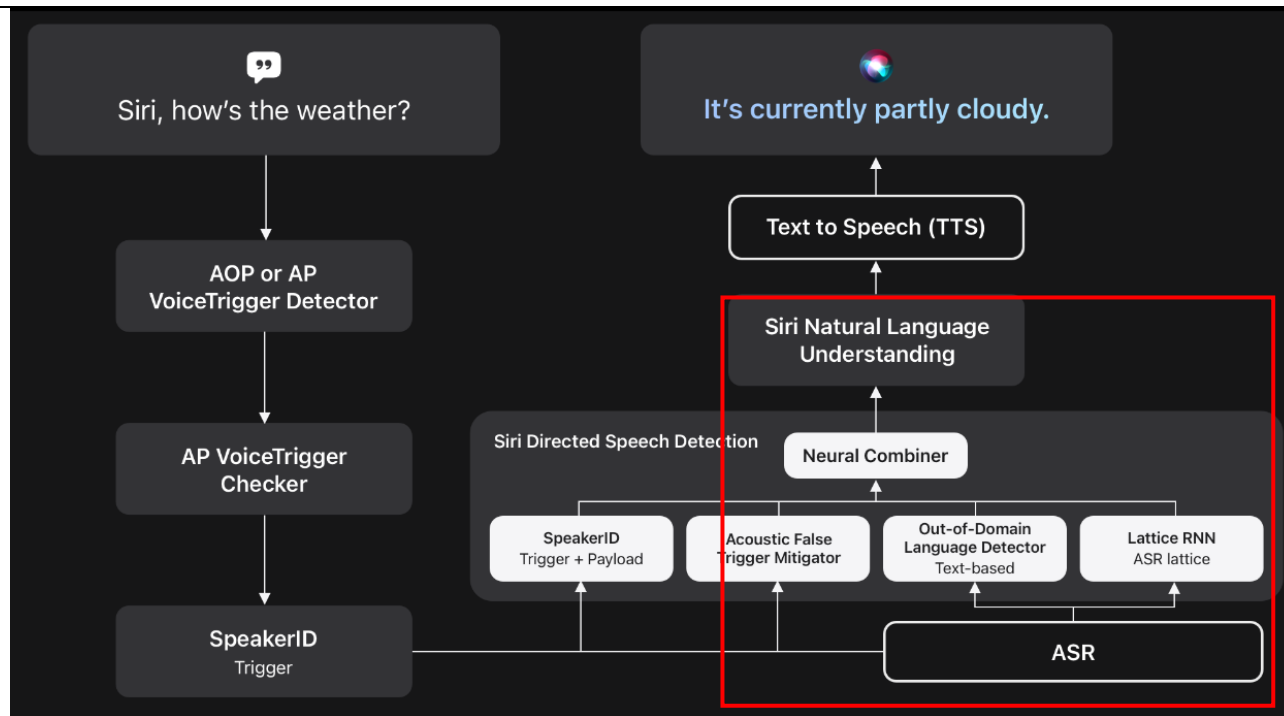
Source: <https://developer.apple.com/design/human-interface-guidelines/siri>

Further, to the extent this element is performed at least in part by Company's software source code, Plaintiff shall supplement these contentions pursuant to production of such source code by the Company.

Exhibit 2 to the Complaint

<p>[1.5] one or more Automatic Speech Recognition (ASR) grammar codes;</p>	<p>Company performs and/or induces others to perform a step of preparing, using one or more processing devices instructions for accessing the information, the instructions including: one or more Automatic Speech Recognition (ASR) grammar codes.</p> <p>This element is infringed literally, or in the alternative, under the doctrine of equivalents.</p> <p>For example, Siri does the language processing and semantic analysis to convert the requests into intents. During semantic analysis, the audio input is matched against a grammar ("one or more Automatic Speech Recognition (ASR) grammar codes ") to produce a semantic interpretation of the input.</p> <div data-bbox="411 623 1199 941" style="background-color: black; color: white; padding: 10px;"> <p>A closer look at intents</p> <p>When people use Siri to ask questions and perform actions, Siri does the language processing and semantic analysis needed to turn their requests into intents for your app to handle. The exception is the personal phrase that people create to run a shortcut: When people speak the exact phrase, Siri recognizes it without doing additional processing or analysis.</p> </div> <p>Source: https://developer.apple.com/design/human-interface-guidelines/siri/</p> <p>1.4 Semantic Interpretation</p> <div data-bbox="422 1125 1738 1279" style="border: 2px solid red; padding: 5px;"> <p>A speech recognizer is capable of matching audio input against a grammar to produce a <i>raw text</i> transcription (also known as <i>literal text</i>) of the detected input. A recognizer may be capable of, but is not required to, perform subsequent processing of the raw text to produce a <i>semantic interpretation</i> of the input.</p> </div> <p>Source: https://www.w3.org/TR/2004/REC-speech-grammar-20040316/#S1.3</p>
--	--

Exhibit 2 to the Complaint



Source: <https://machinelearning.apple.com/research/voice-trigger>

Further, to the extent this element is performed at least in part by Company's software source code, Plaintiff shall supplement these contentions pursuant to production of such source code by the Company.

[1.6] one or more short text string matching codes; and

Company performs and/or induces others to perform a step of preparing, using one or more processing devices instructions for accessing the information, the instructions including: one or more short text string matching codes.

This element is infringed literally, or in the alternative, under the doctrine of equivalents.

Exhibit 2 to the Complaint

For example, Siri does the natural language processing and semantic analysis to convert the requests into the intents and provides string to match against the data. Since, the relevant information is retrieved according to the intent, upon information and belief, the instructions comprise one or more short text string matching codes.

A closer look at intents

When people use Siri to ask questions and perform actions, Siri does the language processing and semantic analysis needed to turn their requests into intents for your app to handle. The exception is the personal phrase that people create to run a shortcut: When people speak the exact phrase, Siri recognizes it without doing additional processing or analysis.

Source: <https://developer.apple.com/design/human-interface-guidelines/siri/>

Overview

Your app likely defines a number of custom types that model the data the app creates or consumes. For example, a music app might define types that represent artists, albums, and tracks. Because those types are unique to your app, the framework can't interpret them until you expose them to system services such as Siri and the Shortcuts app. *Entities* are lightweight types that provide information to the system about your app's data or concepts relating to that data. An entity identifies and queries the data it represents and describes how the system displays that data onscreen.

Source: <https://developer.apple.com/documentation/appintents/integrating-custom-types-into-your-intents>

Exhibit 2 to the Complaint

	<div data-bbox="411 250 1346 505" style="border: 2px solid red; background-color: black; color: white; padding: 10px;"> <p>To let people use arbitrary text to find specific entities, adopt the <code>Entity StringQuery</code> protocol instead. Queries that adopt this protocol cause the system to display a search field above the list of suggested entities. Implement the required <code>entities(matching:)</code> function, and use the provided string to match against your data. For example, a music app might let people search for a specific album by matching against the album name.</p> </div> <p>Source: https://developer.apple.com/documentation/appintents/integrating-custom-types-into-your-intents</p> <p>Further, to the extent this element is performed at least in part by Company's software source code, Plaintiff shall supplement these contentions pursuant to production of such source code by the Company.</p>
<p>[1.7] one or more information formatting codes operative to format a consumer device display; and</p>	<p>Company performs and/or induces others to perform a step of preparing, using one or more processing devices instructions for accessing the information, the instructions including: one or more information formatting codes operative to format a consumer device display.</p> <p>This element is infringed literally, or in the alternative, under the doctrine of equivalents.</p> <p>For example, the intents describe how the system displays (“information formatting codes operative to format a consumer device display”) the data such as dates, times, and addresses.</p>

Exhibit 2 to the Complaint

	<p>Overview</p> <p>Your app likely defines a number of custom types that model the data the app creates or consumes. For example, a music app might define types that represent artists, albums, and tracks. Because those types are unique to your app, the framework can't interpret them until you expose them to system services such as Siri and the Shortcuts app. <i>Entities</i> are lightweight types that provide information to the system about your app's data or concepts relating to that data. An entity identifies and queries the data it represents and describes how the system displays that data onscreen.</p> <p>Source: https://developer.apple.com/documentation/appintents/integrating-custom-types-into-your-intents</p> <p>Siri displays entities like dates, times, addresses and currency amounts in a nicely formatted way. This is the result of the application of a process called inverse text normalization (ITN) to the output of a core speech recognition component. To understand the important role ITN plays, consider that, without it, Siri would display "October twenty third twenty sixteen" instead of "October 23, 2016". In this work, we</p> <p>Source: https://machinelearning.apple.com/research/inverse-text-normal</p> <p>Further, to the extent this element is performed at least in part by Company's software source code, Plaintiff shall supplement these contentions pursuant to production of such source code by the Company.</p>
[1.8] communicating the prepared instructions.	<p>Company performs and/or induces others to perform a step of communicating the prepared instructions.</p> <p>This element is infringed literally, or in the alternative, under the doctrine of equivalents.</p>

Exhibit 2 to the Complaint

For example, Siri communicates the instructions to execute intents to the apps in the user's iPhone such that the relevant information is accessed.

A closer look at intents

When people use Siri to ask questions and perform actions, Siri does the language processing and semantic analysis needed to turn their requests into intents for your app to handle. The exception is the personal phrase that people create to run a shortcut: When people speak the exact phrase, Siri recognizes it without doing additional processing or analysis.

Source: <https://developer.apple.com/design/human-interface-guidelines/siri/>

Exhibit 2 to the Complaint

SiriKit provides the following intents.

Domain (link to developer guidance)	Intents
VoIP Calling	Initiate calls.
Workouts	Start, pause, resume, end, and cancel workouts.
Lists and Notes	Create notes.
	Search for notes.
	Create reminders based on a date, time, or location.
Media	Search for and play media content, such as video, music, audiobooks, and podcasts.
	Like or dislike items.
	Add items to a library or playlist.

Source: <https://developer.apple.com/design/human-interface-guidelines/siri>

Further, to the extent this element is performed at least in part by Company's software source code, Plaintiff shall supplement these contentions pursuant to production of such source code by the Company.

Exhibit 2 to the Complaint

2. List of References

1. <https://support.apple.com/en-us/105020>, last accessed on 19 June, 2024.
2. <https://machinelearning.apple.com/research/voice-trigger#figure1>, last accessed on 19 June, 2024.
3. <https://www.apple.com/siri/>, last accessed on 19 June, 2024.
4. <https://machinelearning.apple.com/research/voice-trigger>, last accessed on 19 June, 2024.
5. <https://machinelearning.apple.com/research/hey-siri>, last accessed on 19 June, 2024.
6. <https://arxiv.org/pdf/2005.00119.pdf>, last accessed on 19 June, 2024.
7. <https://developer.apple.com/design/human-interface-guidelines/siri>, last accessed on 19 June, 2024.
8. <https://developer.apple.com/design/human-interface-guidelines/siri#System-intents>, last accessed on 19 June, 2024.
9. <https://www.w3.org/TR/2004/REC-speech-grammar-20040316/#S1.3>, last accessed on 19 June, 2024.
10. <https://developer.apple.com/documentation/appintents/integrating-custom-types-into-your-intents>, last accessed on 19 June, 2024.
11. <https://machinelearning.apple.com/research/inverse-text-normal>, last accessed on 19 June, 2024.
12. <https://www.verizon.com/smartphones/apple/>, last accessed on 19 June, 2024.